

IN THE CLAIMS

1-13. (Cancelled)

14. (Currently Amended) A method of partially replicating configuration information in a distributed database, comprising:

defining a subset of data processing systems within a cluster system as a resource group, wherein a data processing system may be a member of more than one resource group;

defining configuration data for the resource group by instantiating a configuration object containing configuration and status information for a highly available application corresponding to the resource group and having an associated list of data processing systems within the resource group; and

replicating the configuration data only on each data processing system within the resource group.

15. (Original) The method of claim 14, wherein the step of defining a subset of data processing systems within a cluster as a resource group further comprises:

defining a highly available application and each data processing system designated to manage the application as a resource group.

16. (Original) The method of claim 15, wherein the step of defining a highly available application and each data processing system managing the application as a resource group further comprises:

defining a plurality of resource groups for each highly available application within the cluster, each resource group including all data processing systems managing the corresponding application.

17. (Cancelled)

18. (Currently Amended) The method of claim [[17]] 14, wherein the step of replicating the configuration data only on each data processing system within the resource group further comprises:

replicating the configuration object on each data processing system identified in an owners list associated with the configuration object.

19. (Currently Amended) The method of claim [[17]] 14, wherein the step of replicating the configuration data only on each data processing system within the resource group further comprises:

replicating, on a data processing system, a configuration object for each resource group including the data processing system.

20. (Original) The method of claim 14, further comprising:

maintaining, on a data processing system, a configuration object for each resource group including the data processing system and no configuration objects for other resource groups.

21-22. (Cancelled)

23. (Currently Amended) A computer program product in a computer usable medium, comprising:

instructions defining a subset of data processing systems within a network as a resource group, wherein a data processing system may be a member of more than one resource group;

instructions defining configuration data for the resource group comprising instructions instantiating a configuration object containing configuration and status information for a highly available application corresponding to the resource group and having an associated list of data processing systems within the resource group; and

instructions for replicating the configuration data only on each data processing system within the resource group.

24. (Original) The computer program product of claim 23, wherein the instructions defining a highly available application and each data processing system managing the application as a resource group further comprise:

instructions defining a plurality of resource groups for each highly available application within the network, each resource group including all data processing systems managing the corresponding application.

25-26. (Cancelled)

27. (Currently Amended) A cluster multiprocessing system, comprising:

a plurality of data processing systems segregated into a plurality of resource groups, wherein each of the plurality of data processing systems may be a member of more than one resource group;

a plurality of configuration objects each corresponding to a resource group within the plurality of resource groups wherein each of the plurality of configuration objects contains configuration and status information for a highly available application corresponding to the resource group and an associated owners list of data processing systems within the resource group; and

wherein each of the plurality of configuration objects is replicated only on each data processing system within the resource group associated with the configuration object.

28. (Original) The cluster multiprocessor system of Claim 27, wherein a highly available application and each data processing system designated to manage the application is defined as a resource group.

29. (Original) The cluster multiprocessor system of Claim 28, wherein a plurality of resource groups is defined for each highly available application within the cluster, each resource group including all data processing systems managing the corresponding application.

30. (Cancelled)

31. (Currently Amended) The cluster multiprocessor system of Claim [[30]] 27, wherein each of the plurality of configuration objects are replicated on each data processing system identified in an owners list associated with the configuration object.